Time for a Change: The Variance Gamma Model and Option Pricing

Peter Carr

Apollo Hogan

Harvey Stein

Quantitative Finance R&D

Bloomberg LP

January 11, 2007

Revision: 1.28

Abstract

The most widely used option pricing model is the Black-Scholes model. We motivate an alternative option pricing model called the Variance Gamma (VG) model and demonstrate its implementation in Bloomberg.

Header: /bb/cvs/repo/devcvs/FIR/TREQ/VG/Doc/vgbbmag.tex,v 1.28 2007/01/11 18:01:59 hjstein Exp

I Introduction

We all read the news and many of us here at Bloomberg make a living providing the news. But how can you earn a decent living by trading on the news? For many investors and financial intermediaries, the answer is to trade in equity options. Options give their owners the right to buy or sell some stock for a fixed price by a fixed time. Option prices are largely driven by volatility, and volatility is largely driven by news. Thus, news moves stock prices, and the more news that comes in, the more prices move. The more prices move, the more valuable is a given option. Typically speaking, the more news that is anticipated to arrive, the higher the price at which the option trades on the market. In fact, besides its intrinsic value, the main thing that you need to know to value an option is the amount of news that will arrive by the time the option matures. Option traders refer to this premium over intrinsic value as volatility value. Their job is to assess this volatility value of options.

With so much of a stock option's value tied up with news arrival, it's worth taking a look at how news on the underlying stock reaches the market. If you type `NEWS<Go>` into your Bloomberg terminal, the day's headlines are promptly displayed. Scanning the headlines for any particular stock, you may find multiple news items or no news at all. As you monitor the headlines in real time searching for information on the stock, the waiting time from one news flash to another varies quite randomly. When a news item does come in, clicking on the headline can generally tell you whether the news is good or bad, but not what its precise impact on the stock price will be.

With these observations in mind, let's see how the standard option pricing model treats news arrival. All listed (equity) stock options are American-style and the binomial model is probably the most widely used approach for valuing these options. In the binomial model, the waiting time between price changes is constant. In a typical implementation of the model, prices change exactly once per day, always at exactly the same time, with the only source of uncertainty being whether the price moves up or down. So, we see that the binomial model behaves as if news arrives at a constant rate and as if each news item has the same relative impact on the price.

As already indicated, the key to valuing options correctly is to accurately capture the way news arrives. So how can we improve on this toy model of news arrival? Perhaps surprisingly, the key is to look at the binomial model in just the right light. As Richard Feynman said in his Nobel acceptance speech:

> Theories of the known, which are described by different physical ideas may be equivalent in all their predictions and are hence scientifically indistinguishable. However, they are not psychologically identical when trying to move from that base into the unknown. For different views suggest different kinds of modifications which might be made and hence are not equivalent in the hypotheses one generates from them in one's attempt to understand what is not yet understood.

Instead of querying a binomial model as to whether the first news arrival is good or bad, suppose that we instead ask the model how long we have to wait until the first bad news arrives. If $p \in (0, 1)$ is the probability for each time step that good news arrives, then the probability that it takes at least n time steps before the first bad news arrives is given by p^n. This distribution of arrival times is known as the geometric distribution and is characterized among discrete distributions as the only one having the *memoryless property* — the waiting time to the next event of interest is independent of the history.

A problem with using discrete time models to capture news arrivals is that it becomes impossible for the waiting time to the next news event to be less than the length of the time step. The standard solution to this problem is to work with continuous time models arising from discrete time models by letting the length of each time step get arbitrarily short. The continuous time analog of the geometric distribution is the exponential distribution[1], the unique continuous distribution with the memoryless property.

Suppose that we model the waiting times between arrivals of bad news as independent draws from the same exponential distribution with parameter λ_b. Likewise, we model the waiting times to good news arrivals as independent draws from another independent exponential distribution, with parameter λ_g.

[1]When a waiting time is exponentially distributed with parameter $\lambda > 0$, then the probability that the waiting time is at least $t > 0$ is given by $e^{-\lambda t}$.

Recall that in the binomial model, all good or bad news causes the stock price to jump up or down by a fixed percentage amount. To allow for variation in the impact of a given type of news on stock returns, we draw the return reaction from a *gamma* probability distribution with mean μ and variance ν.[2] The main advantage of working with gamma distributions is that the cumulative impact of several independent news arrivals of the same type is also gamma distributed. Suppose that we model the successive up jumps in stock returns arising from good news as independent draws from the same gamma distribution with mean μ_g and variance ν_g. Likewise, we model the successive down jumps in stock returns by drawing from an independent gamma distribution with mean μ_b and variance ν_b. Since the gamma distribution only generates positive outcomes, we negate the draw to generate a negative return.

At this point, the return process is being modeled as the difference of two independent stochastic processes, one accumulating the impact of good news and the other accumulating the impact of bad news. Both of these processes are called compound Poisson processes, where the jump distribution is gamma in both cases.

We have almost finished our modeling of news arrival, with just one more adjustment in order. The standard Black-Scholes model assumes that returns are generated by a Brownian motion. Hence, the stock price is taken as the position of a particle which is never at rest. In any finite time interval, there are an infinite number of moves, all of which are infinitesimally small. As a result, the process jitters, but it does not jump. However, an examination of the behavior of options prices at short maturity argues for both jitters and jumps. As a result, our final model incorporates both.

Compound Poisson processes result in a finite random number of news arrivals in any finite period of time. To induce jitter, we let the two arrival rates approach infinity. We are left with a four parameter model which has an infinite number of jumps in any finite time interval, and whose return process is generated as the difference of two independent gamma processes.

[2] If a return jump J is gamma distributed with mean μ and variance ν, then its probability density function (PDF) is $p_J(j) = C(\mu, \nu) \times j^{\frac{\mu^2}{\nu}-1} \times e^{-\frac{\mu}{\nu}j}$ for $j > 0$, where the required constant is $C(\mu, \nu) \equiv \left(\frac{\mu}{\nu}\right)^{\frac{\mu^2}{\nu}} / \Gamma\left(\frac{\mu^2}{\nu}\right)$ and $\Gamma(\cdot)$ is the gamma function.

To reduce the parameter count to 3, we can set the ratio of the mean to the standard deviation to be the same for the two gamma processes: $\frac{\mu_g}{\sqrt{\nu_g}} = \frac{\mu_b}{\sqrt{\nu_b}}$. The Bloomberg implementation uses the following re-parametrization of the resulting 3-parameter model.[3] First, we let ν denote the common ratio of the variance to the squared mean: $\nu \equiv \frac{\nu_g}{\mu_g^2} = \frac{\nu_b}{\mu_b^2}$. It turns out that the parameter ν governs the excess kurtosis of the stock return distribution. For each gamma process, the ratio of the variance rate ν to the mean rate μ measures the width of the distribution. We introduce a parameter σ which is proportional to the geometric mean of these ratios: $\sigma \equiv \sqrt{\frac{2}{\nu}}\sqrt{\frac{\nu_g}{\mu_g}\frac{\nu_b}{\mu_b}}$. Not surprisingly, the parameter σ governs the standard deviation of the return distribution. Finally, we introduce a parameter θ which is proportional to the difference between these two ratios: $\theta \equiv \frac{1}{\nu}\left[\frac{\nu_g}{\mu_g} - \frac{\nu_b}{\mu_b}\right]$ and which turns out to have the same sign as the skewness of the stock return distribution. The return distribution is fully described by the three parameters σ, θ, and ν. While the three parameters are not literally equal to the second, third, and fourth moments respectively, each parameter is the main determinant of its corresponding moment. The resulting model is a pure jump process, in contrast to much of the existing option pricing literature.

It's time to give this baby a name. Given its genesis, we could call our progeny the "Difference of Two Independent Gamma Processes" model, but that's a long name to go through life with. It turns out that the return distribution that arises from our process results from integrating a Gaussian PDF over its variance parameter. The weight function used in the integration is a gamma PDF. Therefore, this creation was christened the *Variance Gamma* model (henceforth VG) in the seminal 1990 paper by Madan and Seneta[4]. Over 60 papers[4] have since appeared exploring the properties of the model. The model is now described in textbooks by Joshi[3], Schoutens[7], Cont and Tankov[1], and even the venerable Hull[2]. Just as Black-Scholes takes most option traders some time to digest, the VG model has only slowly crept into plain sight. With the new implementation of the Bloomberg `SKEW` function, it is our hope that this new model will enjoy the popularity it deserves.

[3]For a detailed derivation, we refer the reader to Madan, Carr, and Chang[5].
[4]See `IDOC #2021582<GO>` for this bibliography.

II VG on Bloomberg

The VG model is available on the Bloomberg Terminal through the `SKEW` function. `SKEW` allows you to inspect market deviations from Black-Scholes both on a strike basis and on a maturity basis. Data can be viewed with tables, two-dimensional charts and three-dimensional charts. We will illustrate its utility with an example.

`SPX<Index>SKEW<Go>` loads and displays prices for options on the S&P 500. By default it displays the ask prices for call options with the four nearest maturities.[5] We can display option premiums and VG model premiums, but as can be seen in Figure 1, displaying option premiums isn't especially illuminating. They follow the expected curve — the call premiums tends to zero as the strike tends to infinity, and tends to the line *(spot-strike)* as the strike tends to zero. All reasonable models will have this behavior. Differences between models will only be visible by graphing the differences in prices between the models. As can be seen here, the VG model closely matches the quoted premiums.

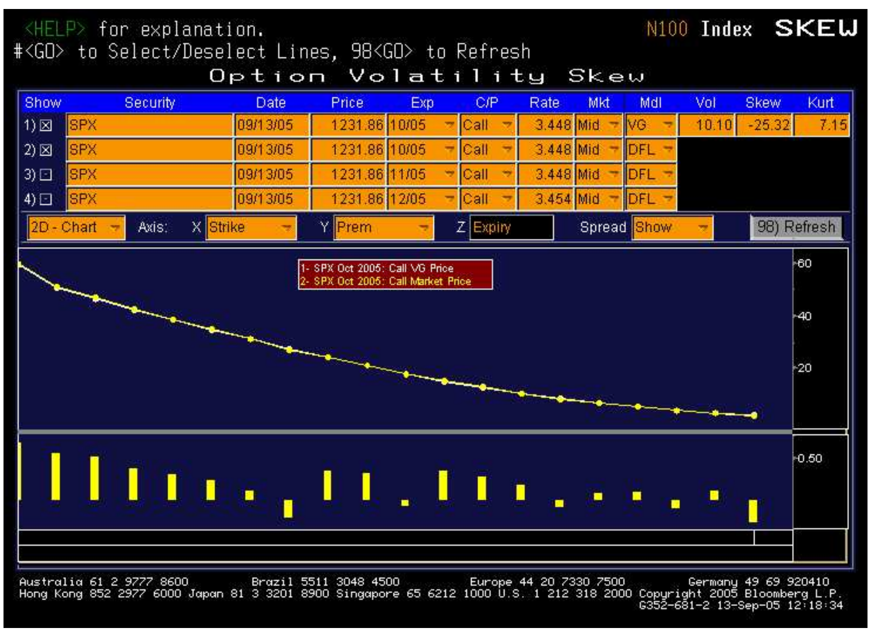

Figure 1: Premium vs. strike for the market and the VG model. The upper graph displays the premiums. The lower graph displays the differences between the VG model prices and the market prices.

[5]For convenience in using `SKEW`, the display settings on the screen are sticky, so that each invocation retrieves the setup you were last using, applying it to the current security.

Displaying the Black-Scholes implied volatility as a function of strike is much more interesting (see Figure 2). Here we see the deviation from the Black-Scholes model typical of the equity and index options markets. (See sidebar.) The volatility at low strikes is higher than at high strikes. We see in this case that the VG model's implied volatilities differ from market implied volatilities by less than 0.6.

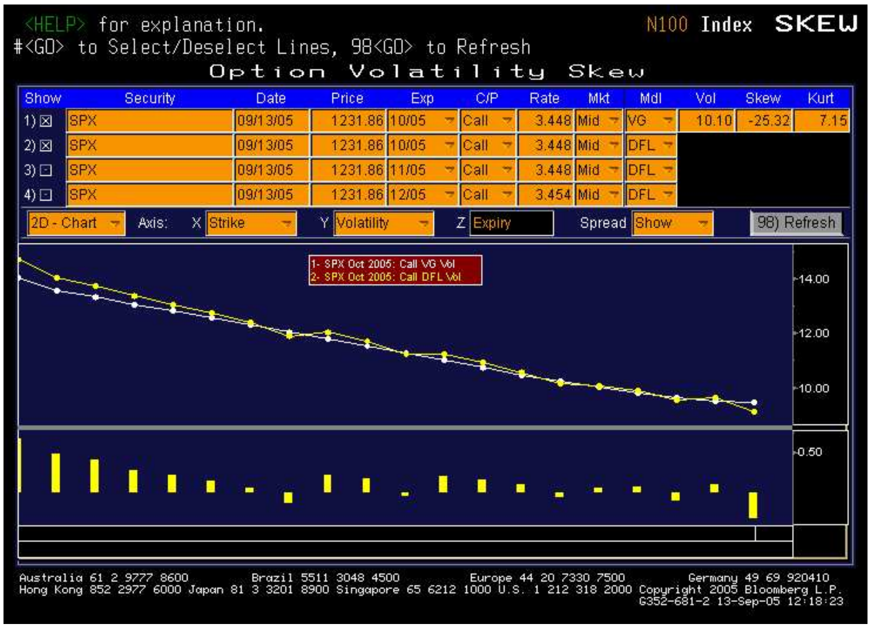

Figure 2: Implied volatility vs. strike for the VG model and for market quotes.

The deviations in implied volatility between the VG implied volatilities and the market implied volatilities can be attributed to a number of factors. One point is that we try to capture the entire market at each maturity. This means fitting the VG model simultaneously to both calls and puts. However, as is seen in Figure 3, despite put-call parity, there are discrepancies in how puts and calls trade. Calls tend to trade at higher implied volatilities than puts.[6] Additionally, out-of-the-money options tend to be more liquid than in-the-money options, further distorting prices when both are used. A third factor is that prices of options at extreme strikes deviate substantially from their theoretical prices because of transaction costs. This introduces further deviations in the model fitting in trying to accomodate these deviant pricings in the wings.

[6]See, for example, Table 1 in Ofek, Richardson, and Whitelaw[6].

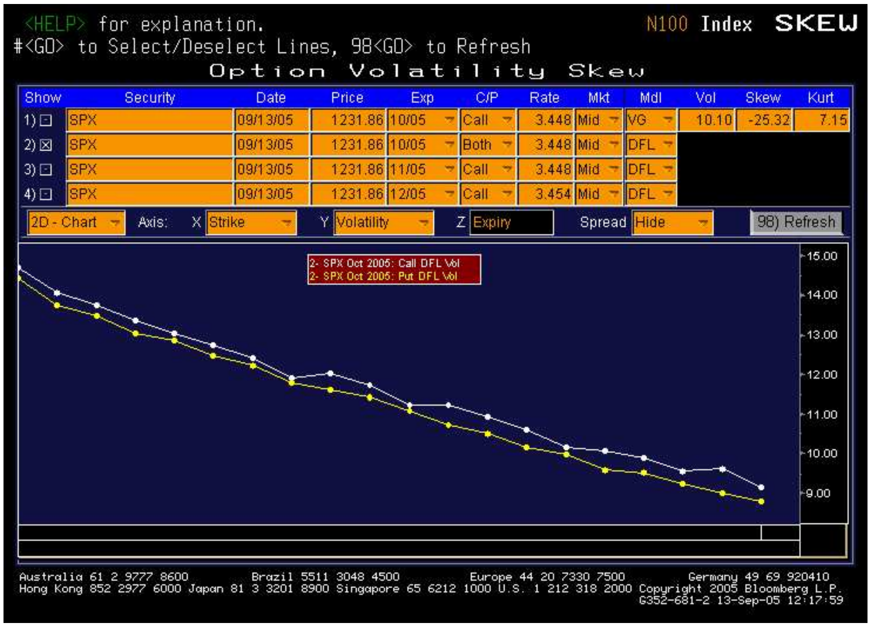

Figure 3: Market implied volatility for calls and puts.

No model is perfect, of course, but in an effort to capture the overall market using the best data available, we fit the VG model to options which are out-of-the-money. This leads to some deviation from the market on the in-the-money sides, and also some torquing of the model at-the-money, as it transitions from call market pricing to put market pricing.

Another useful way of viewing market volatility is to look at all options simultaneously. We can do that with the 3D chart (see Figure 4). Plotting volatility against maturity and strike shows a flattening volatility surface. However, you would expect to see a flattening in strike as maturity increases due to the growth in overall variance over time. To adjust for this, we can plot volatility against maturity and time adjusted moneyness (TAM) $= \ln(X/S_0)/(\sigma\sqrt{t})$, where X is the strike, S_0 is the current spot price, σ is the implied volatility, and t is the time to maturity. TAM expresses the strike in terms of standard deviations from spot instead of in absolute terms. As can be seen in Figure 5, graphing against TAM instead of strike shows that the skew is in fact sustained[7].

Investigating volatility as a function of TAM allows for comparing volatility for different maturities as

[7] Because of its relationship to standard deviation, TAM is presently called "sigma" on the `SKEW` screen.

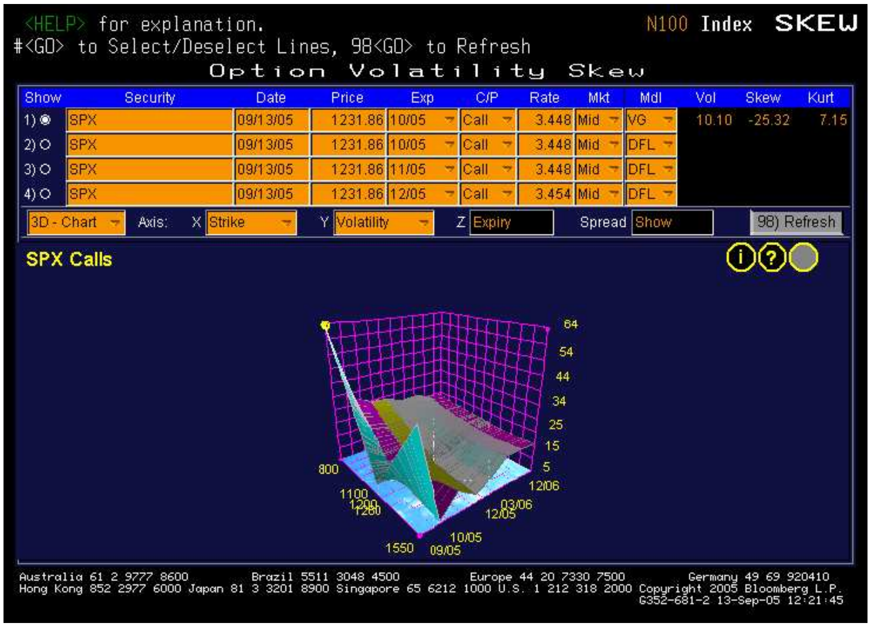

Figure 4: VG implied volatility vs strike and expiry.

well as for different underlyings. You can also adjust for the initial spot price and maturity by plotting with respect to delta, but TAM has the advantage of not compressing the entire range into a finite interval. If you want to only adjust for the initial spot price, you can use moneyness ($= \ln(X/S_0)$) instead of TAM or delta.

Up until now we've worked with mid prices. The reason is that mid pricing is most regular, in that the pricing is uniform in maturity and, by averaging bids and asks, it reduces transaction cost effects. It's also important to look at the prices of actual trades. Comparing the last trade prices of a set of options is difficult. Each trade occurred at a different time, and thus in a different environment. The biggest impact of this is that the spot price was different when each trade occurred. It's also unknown whether each trade is on the bid side or the ask side.

This presents a number of problems. When displaying the implied Black-Scholes volatilities for last trades, each implied volatility has to be computed with the spot price from the time of the trade. Similarly, when fitting the VG model, the fact that the spot price was different for the different trades needs to be taken into account. In the VG implementation, we adjust for this by using the spot price at the time of

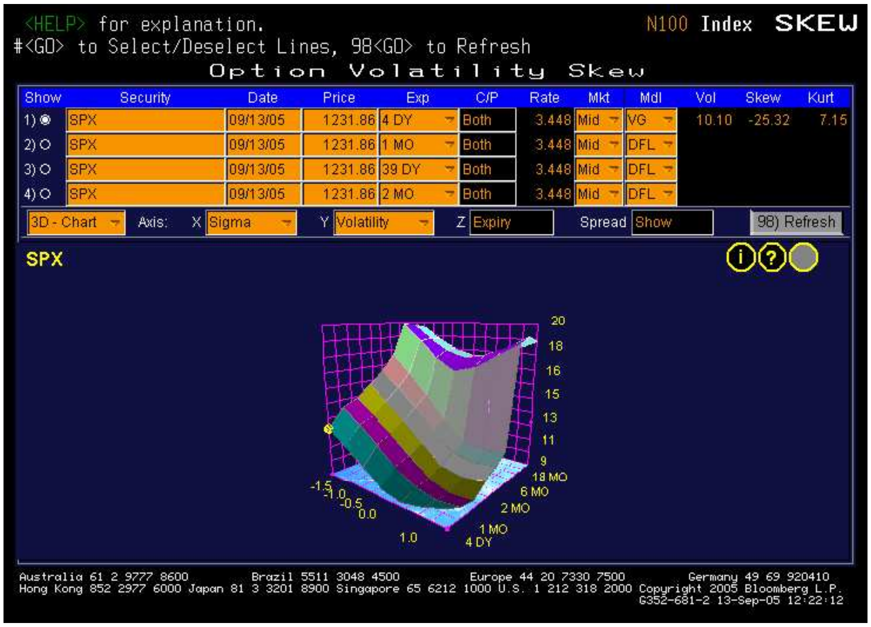

Figure 5: VG implied maturity vs TAM and expiry.

the trade to compute the option's implied volatility.

III Conclusion

Equity option prices are largely driven by stock volatility, and stock volatility is largely driven by the arrival of news. The Black-Scholes model doesn't account for variability in the impact of news or in its arrival rate, leading to a flat volatility smile, which deviates from observed market behavior. The VG model allows for non-deterministic arrival of news, fitting observed market prices much more closely than the Black-Scholes model. The Bloomberg `SKEW` function allows you to explore these relationships.

References

[1] Cont, Rama, and Peter Tankov, 2004, Financial Modelling with Jump Processes, Chapman and Hall.

[2] Hull, John, 2005, Chapter 24, Options, Futures, and Other Derivatives, Sixth Edition, Prentice Hall.

[3] Joshi, Mark, Chapter 18, The Concepts and Practice of Mathematical Finance, Cambridge University Press.

[4] Madan, Dilip, and Eugene Seneta, 1990, "The Variance Gamma (V.G.) Model for Share Market Returns", *Journal of Business*.

[5] Madan, Dilip, Peter Carr, and Eric Chang, 1998, "Variance Gamma Process & Option Pricing", *European Financial Review*, **6**. Also at `http://www.math.ny.edu/research/carrp/research.html`

[6] Ofek, Eli, Matthew Richardson, and Robert Whitelaw, 2004, "Limited arbitrage and short sales restrictions: evidence from the options markets", *Journal of Financial Economics*.

[7] Schoutens, Wim, 2003, Lévy Processes in Finance: Pricing Financial Derivatives, Wiley.

IV Sidebar: Why not just interpolate Black-Scholes?

Developing a model which matches observed option pricing is a complex and difficult task. Even the VG model described here still deviates somewhat from market quotes. This can lead to questioning the effort involved. Instead of developing a sophisticated model (such as the VG model), why not just interpolate and extrapolate prices (or Black-Scholes implied volatilities) directly?

The cost of using such a model is paid back in a number of ways. Most importantly, it gives consistent and arbitrage-free prices for options across strike. Pricing via interpolation gives no guarantee that the computed prices will be arbitrage-free. The problems with extrapolation are worse. How should options with strikes outside of the traded range be priced? Reasonable pricing of these options is important for price quotation and marking to market. Interpolation prices these options by fiat — flat extrapolation implies one price and linear extrapolation another.

Correct computation of Greeks is important for hedging purposes and scenario analysis. Model based Greeks can be used for accurate model based hedging. Using the delta from the Black-Scholes formula ignores the change in option premium arising from the change in implied volatility as the spot price moves. Hedging with this delta will lead to greater hedging errors than necessary.

Finally, typical interpolation schemes yield volatility surfaces that have only a few continuous derivatives. Smoothness is especially important when pricing American options. In the most extreme case, piecewise linear interpolation of implied volatilities implies jumps in the density function. It's unlikely that the market behaves in this fashion.